

9 providing the first work request to a second process on the second node, wherein
10 the first work request specifies that the first process is to directly receive
11 results of the first work;
12 determining based upon the first work and the mapping data, that the first work is
13 also to be performed on a third node from the plurality of nodes, and
14 providing a second work request to a third process on the third node, wherein the
15 second work request specifies that results of the first work performed on
16 the third node are to be provided directly to the first process.

1 16. (ONCE AMENDED) A distributed computing system for performing work, the
2 distributed computing system comprising:
3 a plurality of nodes; and
4 a director communicatively coupled to the plurality of nodes, the director being
5 configured to
6 maintain mapping data that specifies work that can be performed by each
7 of the plurality of nodes,
8 in response to a first work request to perform first work from a first
9 process on a first node from the plurality of nodes, determine based
10 upon the first work and the mapping data, that the first work is to
11 be performed on a second node from the plurality of nodes, and
12 request that the first work be performed by a second process on the second
13 node, wherein the request specifies that first results of the first
14 work be directly provided to the first process;
15 determining based upon the first work and the mapping data, that the first
16 work is also to be performed on a third node from the plurality of
17 nodes, and

18 providing a second work request to a third process on the third node,
19 wherein the second work request specifies that results of the first
20 work performed on the third node are to be provided directly from
21 the third node to the first process.

1 20. (ONCE AMENDED) A computer-readable medium carrying one or more
2 sequences of one or more instructions for processing data on a distributed
3 computing system that includes a plurality of nodes, the one or more sequences of
4 one or more instructions include instructions which, when executed by one or
5 more processors, cause the one or more processors to perform the steps of:
6 maintaining mapping data that specifies work that can be performed by each of
7 the plurality of nodes;
8 in response to receiving a first work request to perform first work from a first
9 process on a first node from the plurality of nodes, determining based
10 upon the first work and the mapping data, that the first work is to be
11 performed on a second node from the plurality of nodes;
12 providing the first work request to a second process on the second node, wherein
13 the first work request specifies that the first process is to directly receive
14 results of the first work;
15 determining based upon the first work and the mapping data, that the first work is
16 also to be performed on a third node from the plurality of nodes, and
17 providing a second work request to a third process on the third node, wherein the
18 second work request specifies that results of the first work performed on
19 the third node are to be provided directly from the third node to the first
20 process.

1 30. (NEW) A method for processing data on a distributed computing system that

b4
2 includes a plurality of nodes, the method comprising the steps of:
3 maintaining mapping data that specifies work that can be performed by each of the
4 plurality of nodes; and
5 in response to receiving a first work request to perform first work from a first
6 process on a first node from the plurality of nodes,
7 determining based upon the first work and the mapping data, that the first
8 work is to be performed on a second node from the plurality of
9 nodes,
10 generating an updated first work request that specifies that the first process
11 is to directly receive results of performing the first work, and
12 providing the updated first work request to a second process on the second
13 node.

1 31. (NEW) An apparatus for processing data on a distributed computing system, the
2 apparatus comprising a memory carrying one or more sequences of one or more
3 instructions which, when executed by one or more processors, cause the one or more
4 processors to perform the steps of:
5 maintaining mapping data that specifies work that can be performed by each of the
6 plurality of nodes; and
7 in response to receiving a first work request to perform first work from a first
8 process on a first node from the plurality of nodes,
9 determining based upon the first work and the mapping data, that the first
10 work is to be performed on a second node from the plurality of
11 nodes,
12 generating an updated first work request that specifies that the first process
13 is to directly receive results of performing the first work, and

14 providing the updated first work request to a second process on the second
15 node.

64
1 32. (NEW) A computer-readable medium for processing data on a distributed
2 computing system, the computer-readable medium carrying one or more sequences
3 of one or more instructions which, when executed by one or more processors, cause
4 the one or more processors to perform the steps of:
5 maintaining mapping data that specifies work that can be performed by each of the
6 plurality of nodes; and
7 in response to receiving a first work request to perform first work from a first
8 process on a first node from the plurality of nodes,
9 determining based upon the first work and the mapping data, that the first
10 work is to be performed on a second node from the plurality of
11 nodes,
12 generating an updated first work request that specifies that the first process
13 is to directly receive results of performing the first work, and
14 providing the updated first work request to a second process on the second
15 node.
